

Submission in Response to NSF CI 2030 Request for Information

DATE AND TIME: 2017-04-04 21:21:50

PAGE 1

REFERENCE NO: 218

This contribution was submitted to the National Science Foundation as part of the NSF CI 2030 planning activity through an NSF Request for Information, https://www.nsf.gov/publications/pub_summ.jsp?ods_key=nsf17031. Consideration of this contribution in NSF's planning process and any NSF-provided public accessibility of this document does not constitute approval of the content by NSF or the US Government. The opinions and views expressed herein are those of the author(s) and do not necessarily reflect those of the NSF or the US Government. The content of this submission is protected by the Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 International License (<https://creativecommons.org/licenses/by-nc-nd/4.0/legalcode>).

Author Names & Affiliations

- Craig A. Stewart - Indiana University
- David Y. Hancock - Indiana University
- Matthew R. Link - Indiana University

Contact Email Address (for NSF use only)

(Hidden)

Research Domain, discipline, and sub-discipline

Computer Science

Title of Submission

Addressing the national need for more diversity and aggregate investment in cyberinfrastructure supporting open science and engineering research

Abstract (maximum ~200 words).

There are tremendous unmet demands for interactive, user-driven advanced cyberinfrastructure tools best delivered from cloud environments. As leaders of the NSF-funded Jetstream cloud system, we set out some of the many pressing scientific needs for cloud-based computing in disciplines ranging from genomics to engineering to field biology to observational astronomy. At the same time, there is increasing demand for traditional resources at the tens of PetaOPs level as well as exascale computing. Going forward, one can expect that NSF investments in cyberinfrastructure will represent a decreasing fraction of total national investment in advanced CI supporting open science and engineering research. There is only one solution to maintaining US global leadership in innovation while meeting the large and diverse CI needs faced by the national research community: a net increase in financial investment in CI resources from a diversity of financial resources. The NSF has within its means and power the ability to do much to coordinate diverse sources of investment and diverse resources that will support open research in academic and government research laboratories throughout the US by providing technical leadership, creating financial incentives, and leveraging investments that include but go far beyond NSF's direct investments in advanced cyberinfrastructure resources.

Question 1 Research Challenge(s) (maximum ~1200 words): Describe current or emerging science or engineering research challenge(s), providing context in terms of recent research activities and standing questions in the field.

The current NSF cyberinfrastructure successfully serves thousands of users who need supercomputers or large clusters advancing critical areas of science and engineering. However, it has only recently started to address the needs of tens of thousands of other NSF-supported

Submission in Response to NSF CI 2030 Request for Information

DATE AND TIME: 2017-04-04 21:21:50

PAGE 2

REFERENCE NO: 218

researchers who do not fit the traditional high performance computing (HPC) model. These are not necessarily researchers who could be served simply by a different scale of HPC, but rather they come from diverse user communities that have only recently come to use large-scale computation as a research tool and don't fit the "standard" model of a uniform software environment in terms of applications, languages, and libraries. Furthermore, the standard operational models for HPC systems do not match the needs of many of these researchers. They often have challenges in common: they have large quantities of digital data requiring multiple uniquely configured computing environments, they need support for reproducible computation, they need environments with support for collaboration, they need storage for their data, and they need to be able to depend upon federally-funded HPC resources. Such researchers not served by traditional HPC resource may often fit into two other categories: they may well understand their research needs but may not easily understand advanced computing well enough to state their needs as clearly as a user of traditional HPC systems; and, the advanced cyberinfrastructure systems most appropriate to meet their needs may not be present at their home institution.

We operate the Jetstream NSF-funded cloud computing system and, in proposing and implementing it, we have found several needs for cloud computers, representing many thousands of users across the US. These needs are described in brief below:

In Genomics, there are multiple large user communities that operate sizable software infrastructure projects. Major NSF-supported software projects include CyVerse and Galaxy. Together the US-based users of these platforms total in the tens of thousands. The experimental biologists who use iPlant and Galaxy use a number of tools, and bioinformatics software is now a tool like any other in the lab, which scientists would prefer to use interactively. Furthermore creators of software tools use Galaxy and iPlant as ways to distribute software and make it easily accessible to practicing biologists.

In Neuroscience, there have been recent efforts in the US (Human Connectome Project) and abroad (UK Brain Bank) to systematize the collection of data from human brains. Data collection is ongoing and has already surpassed the petabyte mark. Given that there are tens of thousands of neuroscientists who would like to access the data -- it is a considerable challenge to provide access to the data and also to the tools to analyze the data. Like many fields of science, Neuroscience currently does not have established mechanisms to ensure the reproducibility of scientific results nor the tools and workflows to leverage traditional HPC environments. Specifically, the sharing of (open) software and various computing resources has been a challenge to this field with the diversity of interests in brain data.

In Earth Sciences, there are large consortia and projects such as UNAVCO and IRIS focused on geodesy and seismology respectively with wide-ranging infrastructure needs than cannot be met by single CI systems. Hosting interactive science gateways with access to petascale storage systems for analysis and storage is the current norm with more interactive analysis through auto-scaling cloud workflows coming in the near future.

In Engineering, tools such as Matlab are important but only a small portion of the story for a domain that is in need of resources to properly scale and leverage traditional CI environments. Projects such as Design Safe CI are starting to couple traditional environments with cloud resources such as Jetstream for more interactive workflows and training, both are needed to support storm surge modeling, computational fluid dynamics, seismic simulations, and data analysis.

Field stations worldwide provide longitudinal data that enhance understanding of basic biology and the impact of human population pressures and global climate change on the environment. A key strength of Biological Field Stations are the ability to combine data across long-term studies to facilitate new science, and much published ecological research depends upon research conducted at field stations. However, it's difficult to maintain a record on site of the research conducted or copy of data collected because researchers come and go throughout the year and managing data collected at field stations is increasingly difficult - partly because of the proliferation of field-deployed instruments each producing its own stream of output data.

In Social Sciences, the increase in availability of "born digital" data has particularly affected the field, especially as much of this relevant data has generated with embedded geographical location information. The Odum Institute at the University of North Carolina is an exemplar of emerging needs for social science research. According to Odum Institute leaders, the computing needs of their data customers far exceed the computing resources available and though individual analysis may be modest their regular users exceed 20,000, creating an aggregate need that is both significant and not currently met.

In Observational Astronomy, new telescopes with gigapixel imagers have enabled collection of new, high-quality data about our universe. Surveys such as the Dark Energy Survey (DES) and PanSTARRS have developed a computational pipeline and secured computational resources to address their specific needs. However, individual researchers use all of the observing time on the One Degree Imager (ODI) and 70% of the time on the Dark Energy Camera (DECam) not dedicated to DES, and have no national-scale computational resources for image analysis. More than 1,000 researchers nationwide who want to analyze data from ODI and DECam.

Submission in Response to NSF CI 2030 Request for Information

DATE AND TIME: 2017-04-04 21:21:50

PAGE 3

REFERENCE NO: 218

Science Gateways are an example of tools that are now considered mainstream and often preferred alternatives to command-line execution of analysis programs. Science Gateways are web-based computational, data analysis, and visualization environments that encode tremendous amounts of expertise in executing sophisticated analyses behind a user-friendly interface. Science Gateways also abstract much detailed information about where or how a particular analysis is done, while also providing a highly controlled and secure way for practicing scientists to make use of large scale and advanced cyberinfrastructure resources. Operating a resource such as Jetstream we've seen the wide-ranging need in many of the above disciplines for persistent hosting from IRIS and UNAVCO in earth sciences, Unidata in meteorology, Galaxy in biology, SEAGrid for engineering, NAMDRunner for molecular dynamics, and many others in chemistry, neuroscience, and medicine. It's no longer sufficient to plan for these services after deployment of a CI system, the API needs and storage resources must be a key element in the design phase.

The orchestrated use of workflows is becoming a standard practice with the wide array of orchestration engines available such as Docker Swarm, Apache Mesos, Kubernetes, and OpenStack Heat. These engines allow researchers and science gateways to dynamically scale applications on cloud resources and deploy them with consistent software environments for reproducibility and reliability. This "mode of use" is now a common expectation and need for many scientists who access current generation orchestration tools directly, as well as Science Gateway developers who use such tools to manage workflows on behalf of Science Gateway users.

Question 2 Cyberinfrastructure Needed to Address the Research Challenge(s) (maximum ~1200 words): Describe any limitations or absence of existing cyberinfrastructure, and/or specific technical advancements in cyberinfrastructure (e.g. advanced computing, data infrastructure, software infrastructure, applications, networking, cybersecurity), that must be addressed to accomplish the identified research challenge(s).

The previous are illustrations of the large number of researchers who are in need of access to advanced computations resources, and certainly not an exhaustive list. These are diverse communities with diverse needs; yet, they share common problems. They use software as a research tool, much of this software needs to be used interactively, they need uniquely configured computing environments, they need reproducible computation, they need environments with support for collaboration, they need storage for their data, and they need these resources to be made available to them in a way that is sustained and which can be depended upon when developing CI plans for NSF-funded research projects.

In terms of technical needs, the sort of resources that these researchers need can be considered programmable cyberinfrastructure, where an evolving API-driven software ecosystem is coupled with commodity hardware, such as what is provided by the Jetstream.

Programmable cyberinfrastructure includes the following sorts of cloud-based characteristics and underlying resources:

- First-class support for virtual machines and container technologies
- High availability environments that keep pace with rapidly evolving software projects
- Integration with the latest orchestration engines
- Large traditional parallel file systems integrated with object storage technologies that are more appropriate for data distribution, searching, and small-file IO
- Larger memory per core
- Longer data life-cycles that minimize data movement, coupled with longer allocation windows
- Availability of specialized processors (GPUs, many-core processors generally, and FPGAs) available in a cloud environment
- Support, but not exclusive deployment of accelerated environments for application optimization (critical for the rise of deep learning and AI approaches)

Other cloud-based resources, that provide for the needs not met in any existing NSF-funded cloud-like resource for production, have some of the following characteristics:

- Self-service, dynamic deployment of large scale MapReduce, Apache Spark, or similar analytics environments
- IoT hubs for ingestion of billions of files with streaming analytics capabilities. (Field stations are representative of the large IoT surge now coming to CI environments where traditional CI environments do not fit the large numbers of small records and distributed data analysis with increased searchability are paramount.
- Dynamic and diverse database services
- More capacity. For example, more than 7,000 users of CyVerse and Galaxy are accessing the Jetstream cloud system -- submitting already more than 100,000+ jobs within the first year of production availability of Jetstream.

Submission in Response to NSF CI 2030 Request for Information

DATE AND TIME: 2017-04-04 21:21:50

PAGE 4

REFERENCE NO: 218

While we are still in the relatively early days of Jetstream and are still cultivating the Jetstream user community, we recognize that there is a tremendous need for capacity of cloud resources that goes far beyond the capacity of the existing Jetstream system. We experience the great need for more resources at the tens of petaFLOPS and petaOPs at each XSEDE Resource Allocation meeting when requests for this type of resource may exceed available capacity by three or four times. And, of course, we all read of the tremendous need for exascale systems. All of these needs will seem more important and larger in aggregate resource demands the better the NSF and research institutions get at attracting the attention of and putting advanced cyberinfrastructure tools into the hands of researchers beyond the longstanding communities of traditional HPC and supercomputer users.

Beyond the technical needs identified above, there is simply more need for more advanced cyberinfrastructure resources of all sorts, with no end to the increases in needs in sight.

Question 3 Other considerations (maximum ~1200 words, optional): Any other relevant aspects, such as organization, process, learning and workforce development, access, and sustainability, that need to be addressed; or any other issues that NSF should consider.

Presently, the NSF budget for cyberinfrastructure (CI) constitutes a minority of national investment for open scientific and engineering research. However, the entire NSF budget would likely be insufficient to meet the aggregate national needs for CI resources to support open science and engineering research. There is only one solution to maintaining US global leadership in innovation while meeting the large and diverse CI needs faced by the national research community: a net increase in financial investment in CI resources from a diversity of financial resources. Here we will elaborate on three recommendations.

1) NSF should lead national integration and interoperability of cyberinfrastructure resources and support by integrating and creating interoperability of the resources it directly funds

Currently, the NSF does not have interoperability between the various CI facilities funded by the NSF. In general, the NSF makes investments across directorates and in FFRDCs (Federally Funded Research and Development Center) in ways that result in many isolated pockets of CI resources. This is likely inescapable without change to NSF funding and planning processes. For example, the lifespan of a typical FFRDC is longer than the lifespan of most NSF-funded nationally accessible CI resources. To maximize the value of its investments in CI, the NSF should leverage the fruits of its funded research and facilities and insist on interoperability. Many of the technical obstacles in enabling some measure of effective interoperability have already been addressed (see <http://hdl.handle.net/2022/20538>). With very modest amounts of funding to promote interoperability (with Project Aristotle as a wonderful example: <https://federatedcloud.org>) combined with significant effort in portfolio management, the NSF could do for the generalized US open research cyberinfrastructure ecosystem what it once did in setting the foundation for the modern Internet.

2) NSF should adjust its own funding formulae and approaches so as to provide for excellent support of a diverse set of advanced cyberinfrastructure resources over long periods of time

Funding for delivery services and resources. The current NSF model of funding of advanced CI resources primarily acquisitions is outdated and counterproductive. First, many new and emerging CI resources are delivered only as services. This poses a challenge when the funding model is rooted in the past of supercomputer and cluster acquisitions. There is simply no way for a proposer to respond to recent HPC acquisition solicitations with a wide variety of existing cloud-based and remote services because they may not be “acquired” in the same way as HPC hardware systems. Second, the model of “20% of acquisition cost for maintenance and operations plus help from XSEDE” is inadequate to operate a national resource effectively. This model has not worked for the Jetstream project, and it did not work for several other prior awardees of other traditional HPC clusters or other novel CI resources. In the long run, this is a disincentive to institutions aiming to operate such resources.

Accountability and auditability in resource proposals. The NSF currently bans “voluntary matching” and inclusion of quantitative fiscal information in Facilities statements. However, the playing field for large CI facility awards is inherently tilted in favor of large institutions because winning an award requires subsidizing the NSF. The NSF should make a limited exception to its current ban on quantitative financial information in Facilities Statements and should require that proposals include complete and full budgets for CI facilities that identify the actual total cost of the facility and the sources of funding of any differences between the NSF budget and the total cost of the facility. This will improve effectiveness of grant evaluation, understanding of the aggregate investment in CI facilities being driven by NSF awards, and the ability of the NSF to evaluate and audit such awards.

Submission in Response to NSF CI 2030 Request for Information

DATE AND TIME: 2017-04-04 21:21:50

PAGE 5

REFERENCE NO: 218

Timescales of awards. Presently, it is difficult for NSF awardees to write a grant proposal or make a project execution plan that has dependence upon NSF-funded CI resources such as those supported by XSEDE. This is because XSEDE and the resources it supports are allocated on a year-by-year basis but nearly all grants have longer operational periods. There are certain risks to making fewer CI resource awards, each with longer timelines. But, making longer awards (e.g. 4 or 5 years for an initial award, renewable once for the same time period as an initial award) would create a dependability in resource continuity that could lead to much more use of NSF CI facilities by NSF grant awardees and more effective use of the NSF budget allocated to CI resources overall. Ideally, NSF budgets and plans would include some sort of long-term resource availability plan perhaps forecasting out a decade into the future.

3) NSF should incentivize investment by US research universities and colleges in cyberinfrastructure (integrated as part of a national ecosystem).

Currently, the NSF does not fully incentivize universities to invest in local CI facilities and, in some cases, creates disincentives to local investments. As an example of the disincentives to local investment, verbatim from a campus cyberinfrastructure leader and said at a Birds-of-a-Feather session at a past SC conference: "For me, XSEDE is a way to offload my most expensive local researchers from my budget and on to the NSF budget." Unfortunately, this is a rational response to the current XSEDE-operated allocation processes (which operate based on guidance from the NSF). Simply put, lower campus investment in resources increases (demand and) priority in competing for the chronically oversubscribed resources allocated via XSEDE.

A more significant problem is that in the calculation of indirect (Facilities and Administration) rates, investment in local cyberinfrastructure facilities is included in two categories - Facilities and Administration. Administration is capped at 26%. In our experiences, administrative expenses routinely exceed the 26% cap, meaning that universities do not get credit in the F&A calculations for their investments in local CI support and operations. There is no specific cap on facilities, which includes depreciation on major capital investments such as local CI facilities (clusters, clouds, supercomputers, and storage). Our experience is that, in spite of the lack of a cap on this category of expenses, research universities very often do not get full credit in the final negotiated F&A rates for all of their facilities investments - critically including local major cyberinfrastructure resources. Simply increasing the limits on administrative fees and changing the negotiations on facilities costs will not serve to incentivize campuses to invest more in CI facilities, because an across-the-board increase in F&A rates in absence of a budget increase simply decreases the amount of the NSF budget that goes to supporting the direct projects costs. Instead, the federal government (including NSF) could separate out CI investments of all kinds as its own category, and decrease the total amount of investment included in the calculations for the administration category. This would amount to a decision that the value of investment specifically in CI resources is greater than in some other areas of investment in this category. Isolating CI in support of research as a distinct area, with its own cap within the F&A formula, would incentivize universities to increase investment in CI and thus lead to a national increase in investment in CI facilities.

Consent Statement

- "I hereby agree to give the National Science Foundation (NSF) the right to use this information for the purposes stated above and to display it on a publically available website, consistent with the Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 International License (<https://creativecommons.org/licenses/by-nc-nd/4.0/legalcode>)."